arXiv:1412.6396v1 [cs.LO] 19 Dec 2014

Existential Second-Order Logic Over Graphs: A Complete Complexity-Theoretic Classification

Till Tantau

Institute of Theoretical Computer Science
Universität zu Lübeck, Germany
tantau@tcs.uni-luebeck.de

December 16th, 2014

Abstract

Descriptive complexity theory aims at inferring a problem's computational complexity from the syntactic complexity of its description. A cornerstone of this theory is Fagin's Theorem, by which a graph property is expressible in *existential second-order logic* (ESO logic) if, and only if, it is in NP. A natural question, from the theory's point of view, is which syntactic fragments of ESO logic also still characterize NP. Research on this question has culminated in a dichotomy result by Gottlob, Kolaitis, and Schwentick: for each possible quantifier prefix of an ESO formula, the resulting prefix class either contains an NP-complete problem or is contained in P. However, the exact complexity of the prefix classes inside P remained elusive. In the present paper, we clear up the picture by showing that for each prefix class of ESO logic, its reduction closure under first-order reductions is either FO, L, NL, or NP. For undirected self-loop-free graphs two containment results are especially challenging to prove: containment in L for the prefix $\exists R_1 \cdots \exists R_n \forall x \exists y$ and containment in FO for the prefix $\exists M \forall x \exists y$ for monadic M. The complex argument by Gottlob, Kolaitis, and Schwentick concerning polynomial time needs to be carefully reexamined and either combined with the logspace version of Courcelle's Theorem or directly improved to first-order computations. A different challenge is posed by formulas with the prefix $\exists M \forall x \forall y$, which we show to express special constraint satisfaction problems that lie in L.

1 Introduction

Fagin's Theorem [9] establishes a tight connection between complexity theory and finite model theory: A language lies in NP if, and only if, it is the set of all finite models (coded appropriately as words) of some formula in *existential second-order logic* (ESO logic). This machine-independent characterization of a major complexity class sparked the research area of descriptive complexity theory, which strives to characterize the computational complexity of languages by the syntactic structure of the formulas that can be used to describe them. Nowadays, syntactic logical characterizations have been found for all major complexity classes, see [13] for an overview, although some syntactic extras (like numerical predicates) are often needed for technical reasons.

When looking at subclasses of NP like P, NL, L, or NC^1, one might hope that syntactic restrictions of ESO logic can be used to characterize them; and the most natural way of restricting ESO formulas is to limit the number and types of quantifiers used. All ESO formulas can be rewritten in prenex normal form as $\exists R_1 \cdots \exists R_r \forall x_1 \exists x_2 \cdots \forall x_{n-1} \exists x_n \, \psi$, where the R_i are second-order variables, the x_i are first-order variables, and ψ is quantifier-free. Formulas like $\phi_{\text{3-colorable}} = \exists R \exists G \exists B \forall x \forall y \big(R(x) \vee G(x) \vee B(x) \wedge (E(x,y) \rightarrow$

$\neg(R(x) \wedge R(y)) \wedge \neg(G(x) \wedge G(y)) \wedge \neg(B(x) \wedge B(y)))\big)$, which describes the NP-complete problem 3-COLORABLE, show that we do not need the full power of ESO logic to capture NP-complete problems: the prefix $\exists R \exists G \exists B \forall x \forall y$ suffices. However, do formulas of the form, say, $\exists R \forall x \exists y\, \psi$ also capture all of NP; or do they characterize exactly, say, P? This question lies at the heart of a detailed study by Gottlob, Kolaitis, and Schwentick [11] entitled *Existential Second-Order Logic Over Graphs: Charting the Tractability Frontier,* where the following dichotomy is shown: For each possible syntactic restriction of the quantifier block of ESO formulas, the resulting *prefix class* either contains an NP-complete problem or is contained in P. For instance, it is shown there that all graph problems expressible by formulas of the form $\exists R \forall x \forall y\, \psi$ lie in P, while some problems expressible by formulas of the form $\exists R \forall x \forall y \forall z\, \psi$ are NP-complete. The dichotomy does not, however, settle the question of whether all of P – or at least some interesting subclass thereof like logarithmic space (L) or nondeterministic logarithmic space (NL) – is described by one of the logical fragments.

1.1 Contributions of This Paper

One cannot really hope to show that the prefix class of, say, the quantifier prefix $\exists R \forall x \forall y$ is equal to P since P $\neq$ NP would follow: This syntactically severely restricted prefix class can be shown [6, Proposition 10.6] to be contained in NTIME(n^k) for some constant k and is thus provably different from NP by the time hierarchy theorem. The best one can try to prove are statements like "this prefix class is contained in P and contains a problem complete for P" or, phrased more succinctly, "the reduction closure of this prefix class is P." Our main result, Theorem 1.1, consists of such statements: *For each possible* ESO *prefix class, its reduction closure under first-order reductions is either* FO*,* L*,* NL*, or* NP*.* In particular, no prefix class yields P as its reduction closure (unless, of course, P = NP or NL = P).

It makes a difference which vocabulary we are allowed to use in our formulas and which logical structures we are interested in: Results depend on whether we consider arbitrary graphs, undirected graphs, undirected graphs without self-loops, or just strings. (In this paper, all considered graphs are finite.) The case of strings has been addressed and settled in [6]. In the present paper we consider the same three cases as in [11]: In our vocabulary, we always have just a single binary relational symbol (E), so all models of formulas are graphs. We then differentiate between directed graphs, undirected graphs, and undirected graphs without self-loops (which we call *basic graphs* for brevity). Note that allowing self-loops, whose presence at a vertex x can be tested with the formula $E(x, x)$, is equivalent to considering basic graphs together with an additional monadic input predicate.

To describe the syntactic fragments of ESO logic easily and succinctly, we use the notation of [11]: The uppercase letter E denotes the presence of an existential second-order quantifier, an optional index as in E_2 denotes the arity of the quantifier, and the lowercase letters a and e denote universal and existential first-order quantifiers, respectively. The prefix type of the formula $\phi_{\text{3-colorable}}$ mentioned earlier is $EEEaa$ (or even $E_1E_1E_1aa$ since the predicates are monadic) and we say that $\phi_{\text{3-colorable}}$ *has prefix type* $EEEaa$ (and also $E_1E_1E_1aa$). We will use regular expressions over the alphabet $\{a, e, E, E_1, E_2, E_3, \dots\}$ to denote patterns of prefix types such as E^*aa for "any number of existential second-order quantifiers followed by exactly two universal first-order quantifiers." To define the three kinds of prefix classes that we are interested in, for a formula ϕ let $\text{MODELS}_{\text{directed}}(\phi) = \{G \mid G \text{ is a directed graph and } G \models \phi\}$, $\text{MODELS}_{\text{undirected}}(\phi) = \{G \mid G \text{ is an undirected graph and } G \models \phi\}$, and $\text{MODELS}_{\text{basic}}(\phi) = \{G \mid G \text{ is a basic graph and } G \models \phi\}$. For instance, $\text{MODELS}_{\text{basic}}(\phi_{\text{3-colorable}}) =$ 3-COLORABLE (ignoring coding issues). Next, for a prefix type pattern P, let $\text{FD}_{\text{directed}}(P) = \{\text{MODELS}_{\text{directed}}(\phi) \mid \phi \text{ has a prefix type in } P\}$ and define $\text{FD}_{\text{undirected}}(P)$ and $\text{FD}_{\text{basic}}(P)$ similarly for undirected and basic graphs. "FD" stands for "Fagin-definable" and Fagin's Theorem can be stated succinctly as $\text{FD}_{\text{strings}}(E^*(ae)^*) =$ NP.

As stated earlier, in the context of syntactic fragments of ESO logic it makes sense to consider reduction closures of prefix classes rather than the prefix classes themselves. It

will not matter much which particular kind of reductions we use, as long as they are weak enough. All our reductions will be *first-order reductions* [13], which are first-order queries with access to the bit predicate or, equivalently, functions computable by a logarithmic-time-uniform constant-depth circuit family.[1] Let us write $A \leq_{\text{fo}} B$ if A can be reduced to B using first-order reductions. Let us write $\overline{\text{FD}}_{\text{directed}}(P) = \{A \mid A \leq_{\text{fo}} B \in \text{FD}_{\text{directed}}(P)\}$ for the reduction closure of $\text{FD}_{\text{directed}}(P)$ and define $\overline{\text{FD}}_{\text{undirected}}(P)$ and $\overline{\text{FD}}_{\text{basic}}(P)$ similarly.

Theorem 1.1 (Main Result)**.** *The following table completely classifies all prefix classes of* ESO *logic over basic graphs (upper part) and undirected and directed graphs (lower part):*[2]

If P is at least one of ...	*and at most one of ...,*	*then*
–	$(ae)^*$, E^*e^*a, E_1ae	$\overline{\text{FD}}_{\text{basic}}(P) = \text{FO}$
E_1E_1ae , E_2ae	E^*ae	$\overline{\text{FD}}_{\text{basic}}(P) = \text{L}$
E_1aa	Eaa	$\overline{\text{FD}}_{\text{basic}}(P) = \text{L}$
E_1eaa	E_1e^*aa	$\overline{\text{FD}}_{\text{basic}}(P) = \text{NL}$
E_1aaa, E_1E_1aa, E_2eaa, E_1eae, E_1aee, E_1aea, E_1aae	$E^*(ae)^*$	$\overline{\text{FD}}_{\text{basic}}(P) = \text{NP}$
–	$(ae)^*$, E^*e^*a	$\overline{\text{FD}}_{\text{undirected}}(P) = \overline{\text{FD}}_{\text{directed}}(P) = \text{FO}$
E_1aa	E_1e^*aa, Eaa	$\overline{\text{FD}}_{\text{undirected}}(P) = \overline{\text{FD}}_{\text{directed}}(P) = \text{NL}$
E_1aaa, E_1E_1aa, E_2eaa, E_1ae	$E^*(ae)^*$	$\overline{\text{FD}}_{\text{undirected}}(P) = \overline{\text{FD}}_{\text{directed}}(P) = \text{NP}$

Note that we always have $\overline{\text{FD}}_{\text{undirected}}(P) = \overline{\text{FD}}_{\text{directed}}(P)$, which is not trivial, especially for the prefix E_1aa: On undirected graphs, using only two universally quantified variables, it seems difficult to express "non-symmetric" properties, suggesting $\text{FD}_{\text{undirected}}(E_1aa) \subseteq \text{L}$. However, using a gadget construction, we will show that $\text{FD}_{\text{undirected}}(E_1aa)$ contains an NL-complete problem.

As an application of the theorem, let us use it to prove EVEN-CYCLE $\in$ L, which is the problem of detecting the presence of a cycle[3] of even length in basic graphs B. The complexity of this problem has been researched for a long time, see [12] for a discussion and variants. The idea is to consider the following ESO formulas:

$$\phi_m = \exists C_1 \cdots \exists C_m \forall x \exists y \Big(E(x,y) \wedge \textstyle\bigvee_{i=1}^{m} \big(C_i(x) \wedge C_{(i \bmod m)+1}(y) \wedge \bigwedge_{j \neq i} \neg C_j(x) \big) \Big). \quad (1)$$

They "describe" the following situation: The basic graph can be colored with m different colors so that each vertex x is connected to a "next" vertex y with the "next" color (with color C_1 following C_m). For $m > 2$, it is not hard to see that $B \models \phi_m$ if, and only if, every connected component of B contains a cycle whose length is a multiple of m. Since ϕ_m has quantifier prefix E^*ae and the graphs are basic, the second row concerning basic graphs in Theorem 1.1 tells us that $B \models \phi_m$ can be decided in logarithmic space. The following algorithm now shows EVEN-CYCLE $\in$ L: In a basic input graph B, replace all edges by length-2 paths, then test whether $C \models \phi_4$ holds for some connected component C of B.

1.2 Technical Contributions

The proofs of the statements $\text{FD}_{\text{basic}}(E^*ae) \subseteq \text{L}$ and $\text{FD}_{\text{basic}}(E_1ae) \subseteq \text{FO}$ require a sophisticated technical machinery. In both cases, our proofs follow the ideas of a 35-page proof of $\text{FD}_{\text{basic}}(E^*ae) \subseteq \text{P}$ in [11]. The central observation concerning the first statement is that the *algorithmically* most challenging part in the proof of [11] is the application of Courcelle's Theorem [5] to graphs of bounded tree width. It has been shown in [8] that there is a

[1] As a technicality, since we use first-order reductions with access to the bit predicate, by FO we refer to "first-order logic with access to the bit predicate," which is the same as logarithmic-time-uniform AC^0.

[2] The "interesting" prefixes, where the complexity classes differ between the two parts, are highlighted.

[3] A cycle in an undirected graph must, of course, have length at least 3 and consist of distinct vertices.

logspace version of Courcelle's Theorem, which will allow us to lower the complexity from P to L when the input graphs have bounded tree width. For graphs of unbounded tree width, we will explain how the other polynomial time procedures from the proof of [11] can be reimplemented in logarithmic space.

To prove $\mathrm{FD}_{\mathrm{basic}}(E_1ae) \subseteq \mathrm{FO}$, we need to lower the complexity of the involved algorithms further. The idea is to again follow the ideas from [11] for E_1^*ae. When there is just a single monadic predicate, certain algorithmic aspects of the proof can be simplified so severely that they can actually be expressed in first-order logic. Note, however, that already a second monadic predicate or a single binary predicate makes the complexity jump up to L, that is, $\overline{\mathrm{FD}}_{\mathrm{basic}}(E_1E_1ae) = \overline{\mathrm{FD}}_{\mathrm{basic}}(E_2ae) = \mathrm{L}$.

Concerning the remaining claims from Theorem 1.1 that are not already proved in [11], two cases are noteworthy: Proving that $\mathrm{FD}_{\mathrm{basic}}(E_1eaa)$ contains an NL-complete problem turns out to require a nontrivial gadget construction. Proving $\mathrm{FD}_{\mathrm{basic}}(E_1aa) \subseteq \mathrm{L}$ requires a reformulation of the problems in $\mathrm{FD}_{\mathrm{basic}}(E_1aa)$ as special constraint satisfaction problems and showing that these lie in L.

1.3 Related Work

The study of the expressive power of syntactic fragments of logics dates back decades; the decidability of prefix classes of first-order logic, for instance, has been solved completely in a long sequence of papers, see [2] for an overview. Interestingly, the first-order Ackermann prefix class ae plays a key role in that context and both E_1ae and E^*ae turn out to be the most complicated cases in the context of the present paper, too. The expressive power of monadic second-order logic (MSO logic) has also received a lot of attention, for instance in [3, 5, 7], but emphasis has been on restricted structures rather than on syntactic fragments.

Concerning syntactic fragments of ESO logic, the two papers most closely related to the present paper are [6] by Eiter, Gottlob, and Gurevich and [11] by Gottlob, Kolaitis, and Schwentick. In the first paper, a similar kind of classification is presented as in the present paper, only over *strings* rather than *graphs*. It is shown there that for all prefix patterns P the class $\mathrm{FD}_{\mathrm{strings}}(P)$ is either equal to NP; is not equal to NP but contains an NP-complete problem; is equal to REG; or is a subclass of FO. Interestingly, two classes of special interest are $\mathrm{FD}_{\mathrm{strings}}(E_1^*ae)$ and $\mathrm{FD}_{\mathrm{strings}}(E_1^*aa)$, both of which are the minimal classes equal to REG (by the results of Büchi [3]). In comparison, by the results of the present paper $\overline{\mathrm{FD}}_{\mathrm{basic}}(E_1^*ae) = \overline{\mathrm{FD}}_{\mathrm{basic}}(E_1E_1ae) = \mathrm{L}$, while $\overline{\mathrm{FD}}_{\mathrm{basic}}(E_1ae) = \mathrm{FO}$, and $\overline{\mathrm{FD}}_{\mathrm{basic}}(E_1^*aa) = \overline{\mathrm{FD}}_{\mathrm{basic}}(E_1E_1aa) = \mathrm{NP}$, while $\overline{\mathrm{FD}}_{\mathrm{basic}}(E_1aa) = \mathrm{L}$.

The present paper builds on the paper [11] by Gottlob, Kolaitis, and Schwentick, which contains many of the upper and lower bounds from Theorem 1.1 for the class NP as well as most of the combinatorial and graph-theoretic arguments needed to prove $\mathrm{FD}_{\mathrm{basic}}(E^*ae) \subseteq \mathrm{L}$ and $\mathrm{FD}_{\mathrm{basic}}(E_1ae) \subseteq \mathrm{FO}$. The paper misses, however, the finer classification provided in our Theorem 1.1 and Remark 5.1 of [11] expresses the unclear status of the exact complexity of $\mathrm{FD}_{\mathrm{basic}}(E^*ae)$ at the time of writing, which hinges on a problem called $\mathrm{SATU}(P)$: "Note also that for each P, $\mathrm{SATU}(P)$ is probably not a PTIME-complete set. [...] This is due to the check for bounded treewidth, which is in LOGCFL (cf. Wanke [1994]) but not known to be in NL." The complexity of the check for bounded tree width was settled only later, namely in a paper by Elberfeld, Jakoby, and the author [8], and shown to lie in L. This does not mean, however, that the proof of [11] immediately yields $\mathrm{FD}_{\mathrm{basic}}(E^*ae) \subseteq \mathrm{L}$ since the application of Courcelle's Theorem is but one of several subprocedures in the proof and since a generalization of tree width rather than normal tree width is used.

1.4 Organization of This Paper

To prove Theorem 1.1, we need to prove the lower bounds implicit in the first column of the theorem's table and the upper bounds implicit in the second column. The lower bounds

are proved in Section 2 by presenting reductions from complete problems for L, NL, or NP. The upper bounds are proved in Section 3, where we prove, in order, $\mathrm{FD}_{\mathrm{basic}}(Eaa) \subseteq \mathrm{L}$, $\mathrm{FD}_{\mathrm{basic}}(E^*ae) \subseteq \mathrm{L}$, and $\mathrm{FD}_{\mathrm{basic}}(E_1ae) \subseteq \mathrm{FO}$ using arguments drawn from different areas.

2 Lower Bounds: Hardness for L and NL

For each of the prefix patterns listed in the first column of the table in Theorem 1.1 we now show that their prefix classes contain problems that are hard for L, NL, or NP. The problems from which we reduce are listed in Table 1. As can be seen, we only need to prove new results for a minority of the classes since the NP cases have already been settled in [11].

Table 1: The lower bounds in Theorem 1.1 are proved by showing that the problems in this table, which are complete for the classes in the claims, are either expressible in the fragment or are at least reducible to a problem expressible in the fragment. The problem UNREACH asks whether there is *no* path from s to t in a directed graph. The problems A_2 and A_3 are explained below.

Claim		*Hard problem*	*Proved where*
Lower bounds for basic graphs			
$\overline{\mathrm{FD}}_{\mathrm{basic}}(E_1E_1ae)$	$\supseteq$ L	A_3	Lemma 2.1
$\overline{\mathrm{FD}}_{\mathrm{basic}}(E_2ae)$	$\supseteq$ L	A_2	Lemma 2.2
$\overline{\mathrm{FD}}_{\mathrm{basic}}(E_1aa)$	$\supseteq$ L	2-COLORABLE	[11, Remark 3.1]
$\overline{\mathrm{FD}}_{\mathrm{basic}}(E_1eaa)$	$\supseteq$ NL	UNREACH	Lemma 2.3
$\overline{\mathrm{FD}}_{\mathrm{basic}}(E_1aaa)$	$\supseteq$ NP	POSITIVE-ONE-IN-THREE-3SAT	[11, Theorem 2.2]
$\overline{\mathrm{FD}}_{\mathrm{basic}}(E_1E_1aa)$	$\supseteq$ NP	3-COLORABLE	[11, Theorem 2.3]
$\overline{\mathrm{FD}}_{\mathrm{basic}}(E_2eaa)$	$\supseteq$ NP	3-COLORABLE	[11, Theorem 2.4]
$\overline{\mathrm{FD}}_{\mathrm{basic}}(E_1eae)$	$\supseteq$ NP	3SAT	[11, Theorem 2.5]
$\overline{\mathrm{FD}}_{\mathrm{basic}}(E_1aee)$	$\supseteq$ NP	NOT-ALL-EQUAL-3SAT	[11, Theorem 2.6]
$\overline{\mathrm{FD}}_{\mathrm{basic}}(E_1aea)$	$\supseteq$ NP	POSITIVE-ONE-IN-THREE-3SAT	[11, Theorem 2.7]
$\overline{\mathrm{FD}}_{\mathrm{basic}}(E_1aae)$	$\supseteq$ NP	POSITIVE-ONE-IN-THREE-3SAT	[11, Theorem 2.8]
Remaining lower bounds for undirected and, thereby, also for directed graphs			
$\overline{\mathrm{FD}}_{\mathrm{undirected}}(E_1aa)$	$\supseteq$ NL	UNREACH	Lemma 2.3
$\overline{\mathrm{FD}}_{\mathrm{undirected}}(E_1ae)$	$\supseteq$ NP	3SAT	[11, Theorem 2.1]

The two special languages A_2 and A_3 in the table are defined as follows: For $m \geq 2$ let $A_m = \{G \mid G$ is an undirected graph in which each connected component contains a cycle whose length is a multiple of $m\}$. These languages are all hard for L: In [4, page 388, remarks for problem UFA] it is shown that the reachability problem for graphs consisting of just two undirected trees is complete for L. Since L is trivially closed under complement, testing whether there is *no* path from a vertex u to a vertex v in a graph consisting of two trees is also complete for L, which in turn is the same as asking whether u and v lie in different trees. To reduce this question to A_m, attach cycles of length $2m$ to both u and v. Then all (namely both) components of the resulting graph contain a cycle whose length is a multiple of m if, and only if, u and v lie in different components. (Using a cycle length of $2m$ rather than m ensures that also for $m = 2$ we attach a proper cycle.)

Lemma 2.1. $A_3 \in \mathrm{FD}_{\mathrm{basic}}(E_1E_1ae)$.

Proof. The discussion following the definition of the formula ϕ_3 from equation (1) shows that $\mathrm{MODELS}_{\mathrm{basic}}(\phi_3) = A_3$ holds; but ϕ_3 has the prefix $E_1E_1E_1ae$ rather than E_1E_1ae. However, from ϕ_3 we can easily build an equivalent formula ϕ_2' that only uses two monadic quantifiers: Instead of using one monadic relation for each of the three colors, we can encode three (even four) colors using only two monadic relations: a vertex x has the first color if

$C_1(x) \wedge C_2(x)$, it has the second color if $C_1(x) \wedge \neg C_2(x)$, the third if $\neg C_1(x) \wedge C_2(x)$, and the fourth if $\neg C_1(x) \wedge \neg C_2(x)$. □

Lemma 2.2. $A_2 \in \mathrm{FD}_{\mathrm{basic}}(E_2ae)$.

Proof. Let $\phi = \exists F \forall x \exists y \big(E(x,y) \wedge F(x,y) \wedge \neg F(y,x) \wedge (F(x,x) \leftrightarrow \neg F(y,y))\big)$. Then ϕ has prefix type E_2ae and we claim $A_2 = \mathrm{MODELS}_{\mathrm{basic}}(\phi)$. To see this, first assume that all components in a basic graph B contain a cycle of even length. For a given component, color the vertices on the cycle alternatively white and black. For black vertices x, let $F(x,x)$ hold, while for white vertices x, let $\neg F(x,x)$ hold. Direct the cycle in some way and let $F(x,y)$ hold for any two consecutive vertices x and y (with respect to the orientation). For all vertices x on the cycle we can now choose a vertex y (namely the next vertex on the cycle) such that the quantifier-free part of ϕ is true. To extend the construction to all vertices, repeatedly pick a vertex x not yet colored, but connected by an edge to an already colored vertex y. Assign the opposite color of y to x, set $F(x,x)$ or $\neg F(x,x)$ accordingly, and let $F(x,y)$ hold. The relation F constructed in this way will now witness $B \models \phi$.

For the other direction, let a relation F be given that witnesses $B \models \phi$ and consider any component of B. The formula ϕ chooses for each vertex x a vertex y; let us call this vertex y the *witness* $w(x)$ of x. Clearly, ϕ enforces that there is an edge between x and $w(x)$ in B. Starting at any vertex x in the component under consideration, consider the sequence $x_1 = x$, $x_2 = w(x_1)$, $x_3 = w(x_2)$, and so on. Trivially, $x_i \neq x_{i+1}$ since there are no self-loops in a basic graph, but we also have $x_i \neq x_{i+2}$ since ϕ enforces $\neg F(x_{i+1}, x_i)$, namely for $x = x_i$, and also $F(x_{i+1}, x_{i+2})$, namely for $x = x_{i+1}$. Now, since the graph is finite, the sequence $(x_1, x_2, \dots)$ must run into a cycle and, as we just saw, this cycle must have length at least 3. Finally, the cycle must have even length since $F(x_i, x_i) \leftrightarrow \neg F(x_{i+1}, x_{i+1})$ holds for all vertices x_i on the cycle and, thus, exactly every second vertex on the cycle has a self-loop attached to it by F. □

Lemma 2.3. UNREACH *reduces to a problem in* $\mathrm{FD}_{\mathrm{basic}}(E_1eaa)$ *and also to a problem in* $\mathrm{FD}_{\mathrm{undirected}}(E_1aa)$.

Proof. Since undirected graphs with self-loops are essentially the same as basic graphs with an additional monadic relation (the self-loops allow us to "mark" vertices) and since a single existential first-order quantifier such as the one in E_1eaa also in some sense allows us to single out a set of vertices (those that are connected to it), we temporarily consider the vocabulary (E^2, S^1), instead of our usual vocabulary (E^2). Logical structures are now graphs together with a set of vertices (modeled by S^1). Our objective is to reduce UNREACH to $\mathrm{MODELS}_{\mathrm{basic}}(\phi)$ where ϕ is an (E^2, S^1)-formula of the form $\exists M \forall x \forall y\, \psi$ for monadic M and quantifier-free ψ. Let (G, s, t) be the input for the reduction, where $G = (V, E)$ is a directed graph and $s, t \in V$. We build a new, basic graph $B = (V_B, E_B)$ and a subset S of B's vertices as follows: For each vertex $v \in V$ there will be four vertices in V_B, designated v, $\bar{v}$, v', and $\bar{v}'$. The vertices v' and $\bar{v}'$ will be called the *shadow vertices* of v and $\bar{v}$. The shadow vertices will form the set S. We have the following undirected edges in B, see Figure 1 for an example of the construction:

1. For every vertex $v \in V$ there are the two edges $\{v, \bar{v}\} \in E_B$ and $\{v', \bar{v}'\} \in E_B$ and also the two edges $\{v, v'\} \in E_B$ and $\{\bar{v}, \bar{v}'\} \in E_B$.[4]
2. For every edge $(u, v) \in E$ of the graph G, there is an edge $\{u, v'\} \in E_B$.
3. There are edges $\{\bar{s}, s'\} \in E_B$ and $\{t, \bar{t}'\} \in E_B$.

[4] Using $\{u, v\}$ to indicate an undirected edge between u and v in a basic graph and, in not-so-slight abuse of notation, even writing $\{u, v\} \in E_B$, helps in distinguishing these edges from directed edges in E. Formally, we mean of course $(u, v) \in E_B$ and $(v, u) \in E_B$; and $E_B \subseteq V \times V$ holds.

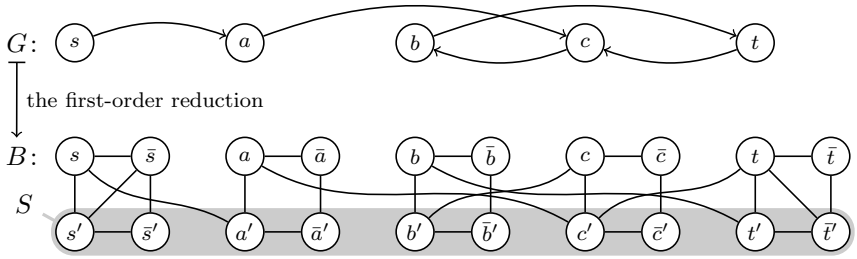

Figure 1: Example of the reduction from Lemma 2.3. The directed graph G on top is reduced to the basic graph at the bottom. The edges from the "squares" are the edges resulting from the first rule, the curved edges result from the second rule, and the two diagonal edges result from the last rule.

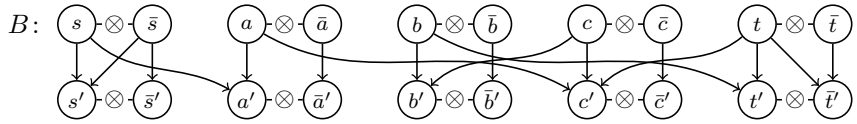

Figure 2: Visualization of the requirements concerning which vertices may lie in M imposed by the formula ψ: For edges with label $\otimes$ exactly one end must lie in M and for directed edges, if the tail of the edge lies in M, the head must also lie in M.

Let ϕ be the following formula:

$$\begin{aligned}\exists M \forall x \forall y \Big(E(x,y) \to \big(\quad & \big((\ S(x) \wedge \ \ S(y)) \to (M(x) \leftrightarrow \neg M(y)) \big) \\ & \wedge \big((\neg S(x) \wedge \neg S(y)) \to (M(x) \leftrightarrow \neg M(y)) \big) \\ & \wedge \big((\neg S(x) \wedge \ \ S(y)) \to (M(x) \to M(y)) \big) \big) \Big).\end{aligned}$$

We make some observations concerning how M can be chosen to make this formula true: First, we only impose restrictions on M when there is an edge between two vertices x and y in B (by "$E(x,y) \to$"). Next, for the edges between vertices inside S ("$S(x) \wedge S(y)$") we require that exactly one of the two endpoints lies in M. The same is true for edges between vertices outside S. Thus, for a vertex v, we always have either $v \in M$ and $\bar{v} \notin M$ or $v \notin M$ and $\bar{v} \in M$. Similarly, we always have either $v' \in M$ and $\bar{v}' \notin M$ or $v' \notin M$ and $\bar{v}' \in M$. The final restriction ("$\neg S(x) \wedge S(y)$") concerns the diagonal and curved edges between a vertex and a shadow vertex: Here, we require that if $x \in M$ holds, we also have $y \in M$. Figure 2 visualizes these restrictions for the example from Figure 1 by placing an $\otimes$-symbol on each edge where exactly one endpoint must be in M and by adding an arrow tip to all edges between a vertex and a shadow vertex.

For any vertex $v \in V$ consider the four vertices v, $\bar{v}$, v', and $\bar{v}'$ in B. Exactly one of v and $\bar{v}$ and exactly one of v' and $\bar{v}'$ must be elements of M. If v is an element of M, then so must v'; and if $\bar{v}$ is an element of M, then so must $\bar{v}'$. This means that a vertex is an element of M if, and only if, its shadow vertex is. Thus, for every vertex $v \in V$ we have $v, v' \in M$ and $\bar{v}, \bar{v}' \notin M$ or we have $v, v' \notin M$ and $\bar{v}, \bar{v}' \in M$. Now consider an edge $(x,y) \in E$. If we have $x \in M$, then we must also have $y' \in M$ and thus, as we just saw, also $y \in M$. This means that when $x \in M$ holds, we also have $z \in M$ for all vertices z reachable from x in G. Now, the edge $\{\bar{s}, s'\}$ in B enforces that $s' \in M$ holds (since one of s and $\bar{s}$ will lie in M and the edge from this vertex to s' enforces that $s' \in M$ holds), which, in turn, enforces $s \in M$. The other way round, the edge $\{t, \bar{t}'\}$ enforces that $t \notin M$ holds since, otherwise, we would have both $t' \in M$ and also $\bar{t}' \in M$, which is forbidden.

Our observations up to now can be summed up as follows: If there is some M that makes ϕ true, there can be *no* path from s to t in G since we must have $s \in M$, $t \notin M$, and together

with s the set M must contain all vertices reachable from s. The other way round, suppose there is no path from s to t in G. Then the formula ϕ is true as the following choice for the set M shows: For each vertex $v \in G$, if v is reachable from s in G, let $v, v' \in M$ and $\bar{v}, \bar{v}' \notin M$; otherwise, let $v, v' \notin M$ and $\bar{v}, \bar{v}' \in M$. Clearly, we now have $s \in M$, $t \notin M$, and all requirements of the formula ϕ are met. This shows that the reduction is correct.

Returning to the original statement of the lemma, we now reduce $\text{MODELS}_{\text{basic}}(\phi)$ to problems in $\text{FD}_{\text{basic}}(E_1 eaa)$ and $\text{FD}_{\text{undirected}}(E_1 aa)$ where there is no S^1-predicate any longer. For this, let ψ be the quantifier-free part of ϕ. We argue that there are (E^2)-formulas ψ' and ψ'' such that $\text{MODELS}_{\text{basic}}(\phi)$ reduces to $\text{MODELS}_{\text{basic}}(\exists M \exists z \forall x \forall y\, \psi')$ and also to $\text{MODELS}_{\text{undirected}}(\exists M \forall x \forall y\, \psi'')$.

Switching over to undirected graphs is fairly easy: Construct ψ'' from ψ by replacing all occurrences of $S(x)$ by $E(x,x)$ and of $S(y)$ by $E(y,y)$. Clearly, we can reduce $\text{MODELS}_{\text{basic}}(\exists M \forall x \forall y\, \psi)$ to $\text{MODELS}_{\text{undirected}}(\exists M \forall x \forall y\, \psi'')$ by mapping a structure (V, E, S) consisting of a basic graph $B = (V, E)$ and a subset $S \subseteq V$ to the undirected graph $(V, E \cup \{(x,x) \mid x \in S\})$.

Next, we wish to replace basic graphs with a designated set S by basic graphs without such a set, but where a special vertex z can be bound by an existential first-order quantifier. Let ψ' be obtained from ψ by replacing all occurrences of $S(x)$ and $S(y)$ by $E(x,z)$ and $E(y,z)$, respectively, and adding the restriction $(x \neq z \wedge y \neq z) \to \ldots$ at the beginning, resulting in the following formula ψ':

$$\begin{aligned}(E(x,y) \wedge x \neq z \wedge y \neq z) \to \big(\quad & \big((\ E(x,z) \wedge \ \ E(y,z)) \to (M(x) \leftrightarrow \neg M(y))\big) \\ & \wedge \big((\neg E(x,z) \wedge \neg E(y,z)) \to (M(x) \leftrightarrow \neg M(y))\big) \\ & \wedge \big((\neg E(x,z) \wedge \ \ E(y,z)) \to (M(x) \to M(y))\big)\big).\end{aligned}$$

We claim that $\text{MODELS}_{\text{basic}}(\exists M \forall x \forall y\, \psi)$ reduces to $\text{MODELS}_{\text{basic}}(\exists M \exists z \forall x \forall y\, \psi')$. The reduction would basically like to map a structure (V, E, S) to a new basic graph B' as follows: B' is identical to $B = (V, E)$, but has a new vertex z^* and edges $\{x, z^*\}$ for all vertices $x \in S$. Then if $(V, E, S) \models \exists M \forall x \forall y\, \psi$, we also have $B' \models \exists M \exists z \forall x \forall y\, \psi'$ since we can choose z^* in $\exists z$. However, the other direction is not clear: It could happen that $B' \models \exists M \exists z \forall x \forall y\, \psi'$, but z is chosen to be some vertex other than z^* and the tests $E(x,z)$, which *should* check whether $S(x)$ used to hold in the original graph, test something different.

To fix this last problem, we modify the construction of B' slightly: We add two triangles p_1, p_2, p_3 and q_1, q_2, q_3 to B' and additionally the two edges $\{z^*, p_3\}$ and $\{z^*, q_3\}$, see Figure 3 for an example. Now, if z is chosen as the vertex z^*, the edges $\{z^*, p_3\}$ and $\{z^*, q_3\}$ mark p_3 and q_3 as shadow vertices and the conditions imposed by ψ' on the triangle can be visualized similarly to Figure 2 as shown also in Figure 3. Clearly, the conditions are satisfied when $p_2, p_3, q_2, q_3 \in M$ and $p_1, q_1 \notin M$.

Now suppose that z is not z^*. We claim that the formula cannot be true in this case: Whatever vertex we choose, the vertices of at least one of the triangles are not connected to the chosen vertex. But, then, ψ' enforces that for each edge of the triangle exactly one end point lies in M, which is not possible in a triangle, yielding a contradiction. □

3 Upper Bounds: Containment in FO and L

The second column of the table in Theorem 1.1 lists upper bounds that we address in the present section. Table 2 shows the order in which we tackle them.

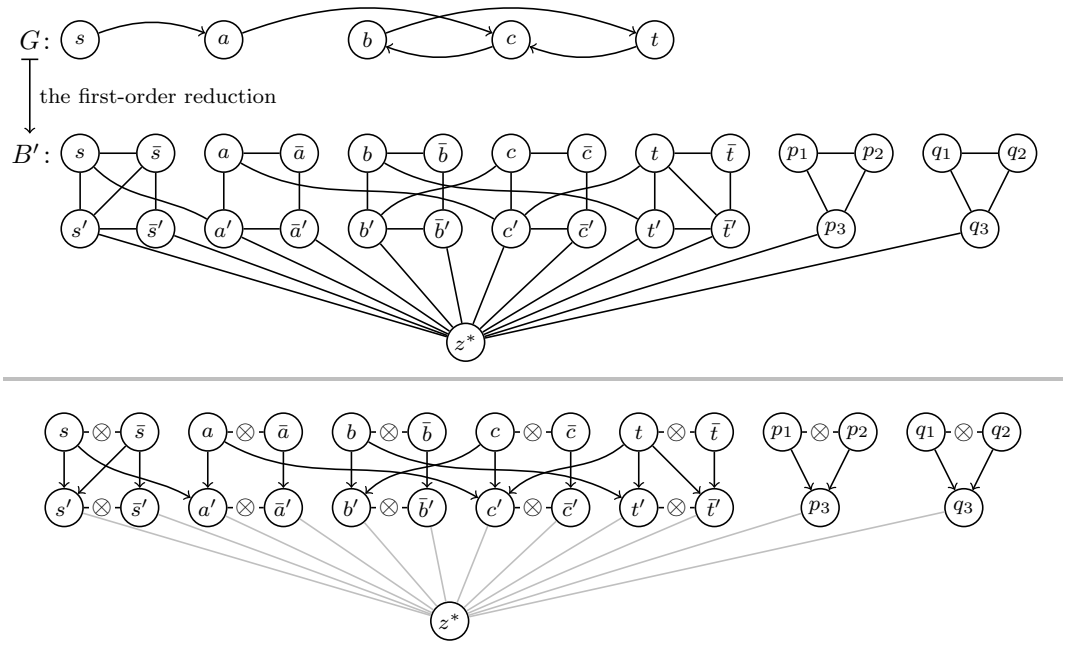

Figure 3: Example of the reduction from UNREACH to $\text{MODELS}_{\text{basic}}(\exists M \exists z \forall x \forall y\, \psi')$ in the upper part. The lower part visualizes the conditions imposed by the formula ψ' when z is chosen to be z^* (nothing is required concerning the gray lines). Note that the conditions on the triangles can easily be satisfied. On the other hand, if any vertex other than z^* is chosen, the conditions in at least one of the triangles will change to three exclusive ors and no solution exists.

3.1 *Eaa* Over Basic Graphs: Reformulation as Constraint Satisfaction

Our first upper bound, $\text{FD}_{\text{basic}}(Eaa) \subseteq \text{L}$, is proved in two steps: First, we reformulate the problems in $\text{FD}_{\text{basic}}(Eaa)$ as special constraint satisfaction problems (CSPs) in Lemma 3.1. Second, we show that these CSPs lie in L in Lemma 3.2.

It will not be necessary to formally introduce the whole theory of constraint satisfaction problems since we will only encounter one very specialized form of them. Furthermore, our CSPs do not quite fit into the standard framework and major results on CSPs like Schaefer's Theorem [15] or the refined version thereof [1] do not settle the complexity of these special CSPs. Nevertheless, we will need some basic terminology: In a binary CSP, we are given a universe U and a set of *constraints*, each of which picks a number of elements from U and specifies one or more possibilities concerning which of these elements may lie in a *solution* $X \subseteq U$. A *constraint language* specifies the types of constraints that we are allowed to use. For instance the constraint language for 3SAT specifies that constraints (which are clauses) must rule out one of the eight possibilities concerning which of the elements (which are the variables) are in X (are set to *true*). We need to deviate from this framework in one important way: we require that there is a constraint for *every* pair of distinct elements of U, not just for some of them. Unfortunately, this deviation inhibits our applying the classification of the complexity of CSPs from [1]; more precisely, the smallest standard CSP classes that are able to express the special CSPs we are interested in are known to contain NL-complete languages – while we wish to prove containment in L.

For sets $C, D \subseteq \{0, 1, 2\}$ we define a $\{C, D\}$-*constraint satisfaction problem* P *on a universe* U to be a mapping that maps each size-2 subset $\{x, y\} \subseteq U$ to either C or D.

Table 2: The upper bounds from Theorem 1.1 and where they are proved. Missing upper bounds for basic and undirected graphs follow from the bounds for directed graphs on the right.

Claims for basic graphs		*Proved where*
$\mathrm{FD}_{\mathrm{basic}}(E_1ae)$	$\subseteq$ FO	Section 3.3
$\mathrm{FD}_{\mathrm{basic}}(E^*ae)$	$\subseteq$ L	Section 3.2
$\mathrm{FD}_{\mathrm{basic}}(Eaa)$	$\subseteq$ L	Section 3.1

Claims for directed graphs		*Proved where*
$\mathrm{FD}_{\mathrm{directed}}((ae)^*)$	$\subseteq$ FO	trivial
$\mathrm{FD}_{\mathrm{directed}}(E^*e^*a)$	$\subseteq$ FO	[11, Theorem 3.1]
$\mathrm{FD}_{\mathrm{directed}}(E_1e^*aa)$	$\subseteq$ NL	[11, Theorem 3.2]
$\mathrm{FD}_{\mathrm{directed}}(Eaa)$	$\subseteq$ NL	[11, Theorem 3.4]
$\mathrm{FD}_{\mathrm{directed}}(E^*(ae)^*)$	$\subseteq$ NP	Fagin's Theorem

A *solution for* P is a subset $X \subseteq U$ such that for all size-2 subsets $\{x, y\} \subseteq U$ we have $|\{x, y\} \cap X| \in P(\{x, y\})$. In other words, P fixes for every pair of two vertices x or y one of two possible constraints concerning *how many* elements of $\{x, y\}$ may lie in X. Let $\textsc{csp}\{C, D\} = \{P \mid P \text{ is a } \{C, D\}\text{-CSP that has a solution}\}$. As an example, $\textsc{csp}\big(\{\{1\}, \{0, 1, 2\}\}\big)$ is essentially the same as the problem 2-COLORABLE = BIPARTITE since a $\{1\}$-constraint enforces that exactly one of two vertices must lie in X (and, hence, corresponds to an edge), while a $\{0, 1, 2\}$-constraint has no effect (and, hence, corresponds to no edge being present). In Lemma 3.2 we show that all $\textsc{csp}\{C, D\}$ lie in L, which is fortunate since we reduce to them:

Lemma 3.1. *For every Eaa-formula ϕ there are sets $C, D \subseteq \{0, 1, 2\}$ such that the set $\textsc{models}_{\mathrm{basic}}(\phi)$ reduces to $\textsc{csp}\{C, D\}$.*

Proof. We may assume that ϕ has the form $\exists M \forall x \forall y\, \psi$ with a *monadic* quantifier M since [11, Lemma 3.3] states that every Eaa-formula is equivalent to an E_1aa-formula. Since the graphs we consider are basic, any occurrence of $E(x, x)$ or $E(y, y)$ in ψ can be replaced by just *false*. Similarly, $E(y, x)$ can be replaced by $E(x, y)$. Finally, we may assume that $\psi \to x \neq y$ holds as well as $\psi(x, y) \leftrightarrow \psi(y, x)$.

Rewrite ψ equivalently as $x \neq y \to \big((E(x, y) \to \gamma) \wedge (\neg E(x, y) \to \delta)\big)$ for formulas γ and δ that are in disjunctive normal form and contain only $M(x)$, $M(y)$, $\neg M(x)$, or $\neg M(y)$ in their terms. Since our graphs are basic and the roles of x and y can be exchanged arbitrarily, γ and δ can only make statements about *how many* elements of the set $\{x, y\}$ lie in M. For instance, if γ is just $M(x)$, then $\forall x \forall y(E(x, y) \to M(x))$ is actually equivalent to $\forall x \forall y(E(x, y) \to (M(x) \wedge M(y)))$ and this imposes the constraint $|\{x, y\} \cap M| = 2$. As further examples, $\gamma = (M(x) \wedge \neg M(y)) \vee (\neg M(x) \wedge M(y))$ imposes the constraint $|\{x, y\} \cap M| = 1$; and $\gamma = M(x) \vee M(y)$ imposes the constraint $|\{x, y\} \cap M| \in \{1, 2\}$. Let C be the cardinality constraints imposed by γ and let D be the cardinality constraints imposed by δ (note that both C and D may be equal to $\emptyset$ or $\{0, 1, 2\}$). Then $\textsc{models}_{\mathrm{basic}}(\phi)$ clearly reduces to $\textsc{csp}\{C, D\}$ by mapping each basic graph B to the following $\{C, D\}$-CSP P: For every edge $\{x, y\}$ of B, let $P(\{x, y\}) = C$; and let $P(\{x, y\}) = D$ when there is no edge $\{x, y\}$ in B. □

Lemma 3.2. *Let $C, D \subseteq \{0, 1, 2\}$. Then $\textsc{csp}\{C, D\} \in$ L.*

Proof. Our aim is to explain, for each choice of C and D, how we can check in logarithmic space whether a $\{C, D\}$-CSP P has a solution $X \subseteq U$. For a given input P, let B be the basic graph whose vertex set is U and which has an edge $\{x, y\}$ when $P(\{x, y\}) = C$. Let $\bar{B}$ be the complement graph of B (exchange edges and non-edges, but do not add self-loops). The edges of B tell us where there are "C-constraints" in P and the edges of $\bar{B}$ where there are "D-constraints" (for $C = D$, the graph $\bar{B}$ is empty, however). We may clearly assume that B has at least three vertices.

We start with some easy observations: If B is the complete graph, then there is always a solution if $0 \in C$ (choose $X = \emptyset$) or $2 \in C$ (choose $X = U$); there is obviously no solution for $C = \emptyset$; and also none for $C = \{1\}$ since the graph contains a triangle while $C = \{1\}$ enforces that B must be bipartite. We can handle $\bar{B}$ being the complete graph similarly. Thus, we may (1) assume that both B and $\bar{B}$ contain at least one edge. This in turn handles

(2) $C = \emptyset$, where there can be no solution, and also none for $D = \emptyset$. On the other hand, (3) if $0 \in C \cap D$ or $2 \in C \cap D$, there is always a solution (namely $X = \emptyset$ or $X = U$). Finally, observe (4) that $\textsc{csp}\{C, D\} = \textsc{csp}\big(\{2 - c \mid c \in C\}, \{2 - d \mid d \in D\}\big)$ since solutions for CSPs of the first kind are the complements of solutions for the second kind.

Let us now go over the cases remaining when $C \neq \emptyset$, $D \neq \emptyset$, $0 \notin C \cap D$, and $2 \notin C \cap D$:

1. $C = \{0\}$. The remaining choices for D are $\{1\}$, $\{2\}$, and $\{1, 2\}$ since otherwise by (3) we are done. For $D = \{1\}$, a solution can only exist if $\bar{B}$ is bipartite and X is one of the shores. Both shores must be non-empty since $\bar{B}$ contains an edge by (1). Since shores are independent sets in $\bar{B}$, the set X must form a clique in B. Since no edge of the clique can satisfy the constraint $C = \{0\}$, there can be no edges and $|X| = 1$. Thus, all we need to check is whether $\bar{B}$ is a star, in which case there will be a solution. Next, for $D = \{2\}$ there can never be a solution since both B and $\bar{B}$ contain an edge, creating conflicting requirements for X. Finally, for $D = \{1, 2\}$ if there is any solution at all, the set $X = \{v \mid v \text{ is isolated in } B\}$ will be such a solution. So, test whether this is indeed the case.

2. $C = \{2\}$. By observation (4) this case is already settled by the previous case.

3. $C = \{0, 2\}$. The only remaining choice for D is $\{1\}$. Again, this means that $\bar{B}$ must be bipartite with shores X and $U \setminus X$. Now, if an edge is missing in $\bar{B}$ between a vertex in X and in $U \setminus X$, the "equality constraint" C cannot be satisfied for this edge in B. Thus, $\bar{B}$ must not only be bipartite, but complete bipartite and, then, there is always a solution. All we need to test is whether $\bar{B}$ is complete bipartite (or, equivalently, whether B consists of two cliques). Clearly, this can be done using even a first-order formula.

4. $C = \{1\}$. The remaining choices are $D = \{1\}$, $D = \{0, 1\}$, $D = \{1, 2\}$, and $D = \{0, 1, 2\}$ (the choices $\{0\}$, $\{2\}$, and $\{0, 2\}$ have already been handled above, with the roles of C and D exchanged). For $D = \{1\} = C$ no solution can exist when the universe has three or more elements, which we assume. For $D = \{0, 1\}$ the situation is similar to the one we had for $C = \{0\}$ and $D = \{1\}$: The constraint $C = \{1\}$ enforces that B is bipartite with one shore being X, but then $D = \{0, 1\}$ enforces that X has size 1. So, again, we just need to test whether a graph is a star, only this time for B. Next, the case $D = \{1, 2\}$ is symmetric to $D = \{0, 1\}$. Finally, for $D = \{0, 1, 2\}$, the only constraint on X is the one given by C, which asks whether B is bipartite. This test can be done in logarithmic space, however, by Reingold's Theorem.

5. $C = \{1, 2\}$. The only remaining choice is $D = \{0, 1\}$. We claim that there is a solution if, and only if, B is a *split graph* (a graph whose vertex set can be partitioned into two sets S_{clique} and S_{indep} such that S_{clique} is a clique and S_{indep} is an independent set). To see this, first note that if B is a split graph, $X = S_{\text{clique}}$ satisfies all constraints: Between vertices inside $X = S_{\text{clique}}$ there are only C-constraint ("pick at least one"), between vertices in $U \setminus X = S_{\text{indep}}$ there are only D-constraint ("pick at most one"), and for every pair of vertices where one lies in X and the other does not, both a C- and a D-constraint is always satisfied. For the other direction, if X is a solution, then there can be no "at most one" constraints between the vertices in X and there can be no "at least one" constraints between the vertices in $U \setminus X$. This shows that X induces a clique in B and $U \setminus X$ induces an independent set in B. Testing whether B is a split graph can be done using a first-order formula since it is known [10] that a graph is a split graph if, and only if, no induced subgraph is isomorphic to $2K_2$, C_4, or C_5.

6. $C = \{0, 1\}$. This is the same as the previous case by observation (4).

7. $C = \{0, 1, 2\}$. No untreated choices for D remain. □

3.2 E^*ae Over Basic Graphs: From P to L

Our objective is to show $\text{FD}_{\text{basic}}(E^*ae) \subseteq \text{L}$ in this section. More precisely, we only need to show $\text{FD}_{\text{basic}}(E_1^*ae) \subseteq \text{L}$ since [11, Theorem 4.1] states $\text{FD}_{\text{basic}}(E^*ae) = \text{FD}_{\text{basic}}(E_1^*ae)$.

A proof of the weaker claim $\text{FD}_{\text{basic}}(E_1^*ae) \subseteq \text{P}$ is spread over the 35 pages of Sections 4, 5, and 6 of the paper [11] by Gottlob et al. and consists of two kinds of arguments: Graph-theoretic and algorithmic. Since the graph-theoretic arguments are independent of complexity-theoretic questions, our main job is to show how the algorithms described by Gottlob et al. can be implemented in logarithmic space rather than polynomial time.

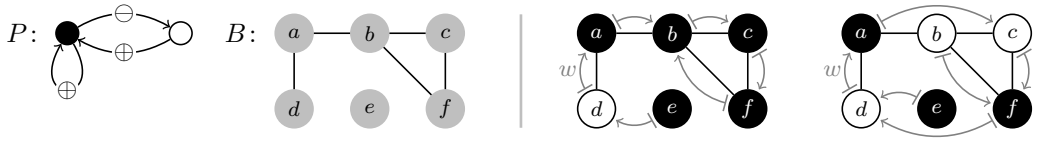

Figure 4: Example of a pattern graph $P = (C, A^{\oplus}, A^{\ominus})$ with two "colors" black and white (so $C = \{black, white\}$, $A^{\oplus} = \{(black, black), (white, black)\}$, and $A^{\ominus} = \{(black, white)\}$) and an uncolored ("gray") example graph B. We have $B \in \text{SATURATION}(P)$ as shown by two examples of legal colorings of B together with witness functions w (in gray).

Similarly to the switch from model checking problems to graphs problems in the previous section, we also wish to reformulate the model checking problems $\text{MODELS}_{\text{basic}}(\phi)$ for E_1^*ae-formulas ϕ in a graph-theoretic manner. Gottlob et al. introduce the notion of *pattern graphs* for this: A *pattern graph* $P = (C, A^{\oplus}, A^{\ominus})$ consists of a set of *colors* C, a set $A^{\oplus} \subseteq C \times C$ of $\oplus$-*arcs*, and a set $A^{\ominus} \subseteq C \times C$ of $\ominus$-*arcs* ($A^{\oplus}$ and $A^{\ominus}$ need not be disjoint). Given a basic graph $B = (V, E)$, a *coloring of* G *with respect to* P is a function $c\colon V \to C$. A mapping $w\colon V \to V$ is called a *witness function for a coloring* c if for all $x \in V$ we have (1) $x \neq w(x)$, (2) if $\{x, w(x)\} \in E$, then $\big(c(x), c(w(x))\big) \in A^{\oplus}$, and (3) if $\{x, w(x)\} \notin E$, then $\big(c(x), c(w(x))\big) \in A^{\ominus}$. If there exists a coloring together with a witness function for B with respect to P, we say that B *can be saturated by* P and the *saturation problem* $\text{SATURATION}(P)$ is the set of all basic graphs that can be saturated by P, see Figure 4 for an example.

The intuition behind these definitions is that a witness function tells us for each x in $\forall x$ which y in $\exists y$ we must pick to make a formula ϕ of the form $\exists M_1 \cdots \exists M_n \forall x \exists y\, \psi$ true. The pattern graph encodes the restrictions imposed by ψ and the monadic predicates M_i:

Fact 3.3 ([11, Theorem 4.6])**.** For every formula $\phi = \exists M_1 \cdots \exists M_n \forall x \exists y\, \psi$, where the M_i are monadic and ψ is quantifier-free, there is a pattern graph P with 2^n vertices such that $\text{MODELS}_{\text{basic}}(\phi) = \text{SATURATION}(P)$.

Thus, it remains to show $\text{SATURATION}(P) \in \text{L}$ for all pattern graphs P. Towards this aim, for a fixed pattern graph P we devise logspace algorithms that work for larger and larger classes of basic graphs B, ending with the class of all basic graphs.

Graphs of Bounded Tree Width and Special Graphs We start by considering only graphs of *bounded tree width,* an important class of graphs introduced by Robertson and Seymour in [14]: A *tree decomposition* of a graph B is a tree T together with a mapping that assigns subsets of B's vertices (called *bags*) to the nodes of T. The bags must have two properties: First, for every edge $\{x, y\}$ of B there must be some bag that contains both x and y. Second, the nodes of T whose bags contain a given vertex x must be connected in T. The *width* of a decomposition is the size of its largest bag (minus 1 for technical reasons). The *tree width* of B is the minimal width of any tree decomposition for it. A class of graphs has *bounded tree width* if there is a constant c such that all graphs in the class

have tree width at most c. From an algorithmic point of view, many problems that can be solved efficiently on trees can also be solved efficiently on graphs of bounded tree width. Courcelle's Theorem turns this into a precise statement:

Fact 3.4 (Courcelle's Theorem, [5])**.** For every MSO-formula ϕ and $t \geq 1$ we have

$$\textsc{models}_{\text{basic}}(\phi) \cap \{G \mid G \text{ has tree width at most } t\} \in \text{LINTIME}.$$

Gottlob et al. apply this theorem to show that when the input graphs B have bounded tree width, we can decide whether $B \in \textsc{saturation}(P)$ holds in polynomial time: the property $B \in \textsc{saturation}(P)$ is easily described in MSO logic. We can lower the complexity from "polynomial time" to "logarithmic space" by using the following logarithmic space version of Courcelle's Theorem:

Fact 3.5 (Logspace Version of Fact 3.4, [8])**.** For every MSO-formula ϕ and $t \geq 1$ we have

$$\textsc{models}_{\text{basic}}(\phi) \cap \{G \mid G \text{ has tree width at most } t\} \in \text{L}.$$

In their graph-theoretic arguments, Gottlob et al. encounter not only graphs of bounded tree width, but also graphs that they call (k,t)-*special* and which are defined as follows: For a basic graph $B = (V, E)$ let us call two vertices u and v *equivalent* if for all $x \in V \setminus \{u, v\}$ we have $\{u, x\} \in E$ if, and only if, $\{v, x\} \in E$. Observe that this defines an easy-to-check equivalence relation on the vertices of B and that each equivalence class is either a clique or an independent set of B. A graph is (k,t)-*special* if we can remove (up to) k equivalence classes $A_1, \ldots, A_k$ from the graph such that the remaining graph has tree width at most t.

The intuition behind (k,t)-special graphs is that equivalent vertices are "more or less indistinguishable" and, thus, for a large enough equivalence class removing some vertices does not change whether the graph can be saturated or not. Formally, let B be (k,t)-special and let $A_1, \ldots, A_k$ be to-be-removed equivalence classes. We obtain an *s-shrink of B* by repeatedly removing vertices from those A_i that have more than s vertices until all of them have at most s vertices. The proof of Lemma 6.4 in [11] implies the following two facts:

Fact 3.6. For every k, t, and pattern graph P there is an s such for every s-shrink B' of a (k,t)-special graph B we have $B \in \textsc{saturation}(P)$ if, and only if, $B' \in \textsc{saturation}(P)$.

Fact 3.7. An s-shrink of a (k,t)-special graph has tree width at most $t + sk$.

In Lemmas 6.3 and 6.4 of [11], Gottlob et al. present polynomial-time algorithms for testing whether a graph is (k,t)-special and for computing an s-shrink when the test is positive. The following lemma shows that we can reimplement these algorithms in a space-efficient manner (which the original algorithms are not):

Lemma 3.8. *For every s, k, and t, there is a logspace computable function that maps every (k,t)-special graph B to an s-shrink of B (and all other graphs to "not (k,t)-special").*

Proof. To check whether a basic graph B is (k,t)-special, simply iterate over all tuples $(v_1, \ldots, v_k)$ of vertices, remove all vertices equivalent to any v_i, and test whether the remaining graph has tree width at most t using the logspace algorithm from Fact 3.5. When a tuple passes the test, for each v_i remove all but the lexicographically first s vertices that are equivalent to v_i from the graph. What remains is the desired shrink. □

The following lemma sums up the bottom line of the above discussion:

Lemma 3.9. *For every pattern graph P and all k and t we have*

$$\textsc{saturation}(P) \cap \{B \mid B \textit{ is } (k,t)\textit{-special}\} \in \text{L}.$$

Proof. Let B be a basic input graph. First, use the algorithm from Lemma 3.8 to (1) test whether B is (k,t)-special (and if not, reject) and then to (2) compute a shrink B' of B. By Fact 3.6 we have $B \in \text{SATURATION}(P)$ if, and only if, $B' \in \text{SATURATION}(P)$. Thus, it suffices to decide the latter membership problem. However, by Fact 3.7 the graph B' has bounded tree width and, thus, we can use the logspace version of Courcelle's Theorem from Fact 3.5 to decide whether $B' \in \text{SATURATION}(P)$ holds. □

Graphs With Self-Saturating Mixed Cycles We extend the class of graphs that our logspace machines can handle to graphs that are not necessarily (k,t)-special, but at least contain a *mixed self-saturating cycle*. A *self-saturating cycle of a basic graph $B = (V,E)$ with respect to a pattern graph $P = (C, A^{\oplus}, A^{\ominus})$* is a sequence $(v_1, v_2, \ldots, v_{n+1})$ of vertices in V for $n \geq 2$ where the v_i for $i \in \{1,\ldots,n\}$ are all different, $v_{n+1} = v_1$, and we can assign colors $c\colon \{v_1,\ldots,v_n\} \to C$ such that for all $i \in \{1,\ldots,n\}$ we have: if $\{v_i, v_{i+1}\} \in E$, then $(c(v_i), c(v_{i+1})) \in A^{\oplus}$; and if $\{v_i, v_{i+1}\} \notin E$, then $(c(v_i), c(v_{i+1})) \in A^{\ominus}$. In other words, B restricted to $\{v_1,\ldots,v_n\}$ can be saturated with the "natural" witness function that "moves along" the cycle. The following is an easy observation concerning self-saturating cycles:

Lemma 3.10. *For every $B \in \text{SATURATION}(P)$ there is a self-saturating cycle in B for P.*

Proof. Let $B = (V,E)$ be saturated with respect to $P = (C, A^{\oplus}, A^{\ominus})$ via some coloring $c\colon V \to C$ and a witness function $w\colon V \to V$. Starting at any vertex v, consider the sequence $v_1 = v$, $v_2 = w(v_1)$, $v_3 = w(v_2)$, …, which must clearly run into a cycle at some point. Let $(v_i, v_{i+1}, \ldots, v_j)$ with $v_j = v_i$ be this cycle. (For instance, in Figure 4 in the first example, starting at e, we run into the cycle (b,c,f,b); and in the second example, starting at e, we run into the cycle (d,a,c,f,d).) Clearly, the cycle $(v_i, v_{i+1}, \ldots, v_j)$ is self-saturating as demonstrated by the coloring c. □

A self-saturating cycle is *mixed* if for some $i,j \in \{1,\ldots,n\}$ we have $\{v_i, v_{i+1}\} \in E$ and $\{v_j, v_{j+1}\} \notin E$, otherwise the cycle is called *pure*. In Figure 4, (b,c,f,b) is a pure self-saturating cycle and (a,c,f,d,a) is a mixed self-saturating cycle as proved by the two example colorings. Two facts concerning mixed self-saturating cycles will be important:

Fact 3.11 ([11, Lemma 6.5])**.** For every pattern graph P there is a constant d such that every basic graph that has a mixed self-saturating cycle with respect to P also has such a cycle of length at most d.

Fact 3.12 ([11, Section 6.3])**.** For each pattern graph P there exist k and t such that $B \in \text{SATURATION}(P)$ holds for all graphs B that contain a mixed self-saturating cycle but are not (k,t)-special.

Lemma 3.13. *For every pattern graph P, we have*

$$\text{SATURATION}(P) \cap \{B \mid B \text{ contains a mixed self-saturating cycle}\} \in \text{L}.$$

Proof. Let k, t, and d be the constants from Facts 3.11 and 3.12. By Fact 3.11, we can decide whether an input graph B contains a mixed self-saturating cycle by iterating over all possible cycles of maximum length d and then testing for all possible colorings whether a saturation has been found for the cycle. If B fails these tests, we can clearly reject.

Otherwise, B has a mixed self-saturating cycle. Test whether B is (k,t)-special using Lemma 3.8 and, if so, use Lemma 3.9 to decide whether $B \in \text{SATURATION}(P)$ holds. Finally, if B is not (k,t)-special, we can accept by Fact 3.12. □

Arbitrary Basic Graphs The last step is to extend our algorithm to graphs that do not contain mixed self-saturating cycles (and are not (k,t)-special, but this will no longer be important). Clearly, by considering the union of the languages from Lemma 3.13 above and Lemma 3.14 below, we see that $\text{SATURATION}(P) \in \text{L}$ holds for all pattern graphs P.

Lemma 3.14. *For every pattern graph P, we have*

$$\text{SATURATION}(P) \cap \{B \mid B \text{ contains no mixed self-saturating cycle}\} \in \text{L}.$$

Proof. Let B be our input graph. Using Fact 3.11 we can first rule out (even using a first-order formula) those B containing a mixed self-saturating cycle. Thus, for $B \in \text{SATURATION}(P)$ to hold, all self-saturating cycles of B must be pure (the reverse is not true, however: B could have a pure self-saturating cycle that cannot be extended to a coloring of the whole graph). In [11], this situation is addressed in Theorem 5.17, which states (reformulated in the terminology of the present paper): *There is a polynomial-time Turing machine that decides* $\text{SATURATION}(P)$ *correctly whenever all self-saturating cycles of the input graph G are pure.* For the proof of this statement, the actual algorithm is summarized at the end of [11, Theorem 5.14] as follows: "In fact, the computationally relevant actions of the algorithm described in this proof are: — Computing the complement G^c of G [...]. — Determining the connected components of G or G^c [...]. — Checking for each component, whether its treewidth is smaller than a constant [...]. — Performing a constant number of further [...] actions on single components, such as the procedure calls $satucheck_P(G)$ or $satucheck'_P(G)$." The omitted parts ("[...]") are statements about the *time complexity* of these operations.

To see that these operations can also be performed in logarithmic space, first note that the complement graph G^c ($\bar{G}$ in the notation of this paper) of G is obtained by simply exchanging edges and non-edges (without introducing self-loops, of course). Determining the connected components of an undirected graph can be done in logarithmic space using Reingold's algorithm. Determining the tree width of a component can be done in logarithmic space [8]. Finally, the procedure calls "$satucheck_P(G)$ or $satucheck'_P(G)$" consist of checking whether a graph G of bounded tree width satisfies a fixed MSO formula, which can be done in logarithmic space by Fact 3.5. □

3.3 E_1ae Over Basic Graphs: From L to FO

Our final task for this paper is showing $\text{FD}_{\text{basic}}(E_1ae) \subseteq \text{FO}$.[5] By Fact 3.3, it suffices to show $\text{SATURATION}(P) \in \text{FO}$ for all pattern graphs with *two* colors (denoted "white" and "black" in the following) and this will be our objective in this section.[6]

In the previous section we proved $\text{SATURATION}(P) \in \text{L}$ for all pattern graphs by developing logspace algorithms that worked for larger and larger classes of graphs. However, this approach is bound to fail for the class FO since properties like "the graph is a tree" (let alone "the graph is (k,t)-special") are not expressible in first-order logic. Instead, in this section we show $\text{SATURATION}(P) \in \text{FO}$ directly for each possible pattern graph with two colors.

The simplest case arises when $P = (C, A^{\oplus}, A^{\ominus})$ is acyclic (meaning that the directed graph $(C, A^{\oplus} \cup A^{\ominus})$ is acyclic): Lemma 3.10 shows that we then have $\text{SATURATION}(P) = \emptyset$ since self-saturating cycles cannot exist for such P. Thus, we only need to consider pattern graphs P with cycles (self-loops are also cycles, here). Since P only has two colors, there are only few ways in which such cycles may arise. The more cycles there are, the easier it will be to color the graph, so we first handle the case that there are cycles both in $A^{\oplus}$ and $A^{\ominus}$, then that there is a cycle in $A^{\oplus}$ or in $A^{\ominus}$, and finally that there is only a cycle in $A^{\oplus} \cup A^{\ominus}$.

[5] In contrast, Lemmas 2.1 and 2.2 show that if we have *two* monadic quantifiers or one *binary* quantifier, the prefix class contains an L-complete problem.

[6] In contrast, using three colors we can describe L-complete problems: $\text{SATURATION}(P) = A_3$ where P contains a $\oplus$-labeled 3-cycle and A_3 is the L-complete language from Table 1.

Lemma 3.15. *Let $P = (\{black, white\}, A^{\oplus}, A^{\ominus})$ contain cycles both in $A^{\oplus}$ and $A^{\ominus}$. Then* SATURATION(P) *contains all graphs with at least two vertices (and is hence in* FO*).*

Proof. Suppose all vertices of B have degree at least 1. Then $B \in$ SATURATION(P) holds for one of two reasons:

1. If there is a self-loop in $A^{\oplus}$ at one of the colors (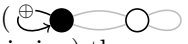 or where the gray arcs can be arbitrary and also be missing) then we can simply color all vertices with the color of the self-loop. The witness function can be set to $w(v) = u$ where u is any neighbor of v.

2. If there is no self-loop in $A^{\oplus}$, the cycle in $A^{\oplus}$ must be 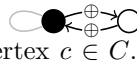 . We treat each connected component C of B separately. Pick any vertex $c \in C$. For each vertex v of the component, color it white if it has an even distance from c, otherwise color it black. Setup the witness function w as follows: Map c to any of its neighbors. Map each vertex v in the component to one of its neighbors that has distance 1 less from c. Clearly, such a neighbor must exist and it will have the opposite color from v.

Now suppose that there is a vertex in B that has degree 0. Then in the complement graph $\bar{B}$ all vertices have an edge to this vertex and, hence, all have degree at least 1. We can now repeat the above argument, only for a cycle in $A^{\ominus}$ instead of $A^{\oplus}$. □

Lemma 3.16. *Let $P = (\{black, white\}, A^{\oplus}, A^{\ominus})$ contain a cycle in $A^{\oplus}$ or in $A^{\ominus}$. Then* SATURATION$(P) \in$ FO*.*

Proof. By possibly switching to complement graphs, we may assume that there is a cycle in $A^{\oplus}$. We may also assume that there is *no* cycle in $A^{\ominus}$ since, otherwise, we can apply Lemma 3.15. As in the proof of that lemma, if in the basic input graph $B = (V, E)$ all vertices have degree at least 1, then $B \in$ SATURATION(P) holds; so assume that there is a vertex of degree 0 in B. Then $A^{\ominus} = \emptyset$ implies $B \notin$ SATURATION(P) since there cannot be an edge between a degree-0 vertex and its witness. Similarly, if all vertices of B have degree 0, then $B \notin$ SATURATION(P): Since $A^{\ominus}$ is acyclic, there is no way to assign a color to all vertices. So, in the following we may assume that the set $S = \{v \mid v \text{ has degree at least 1 in } B\}$ is neither empty nor all of V and that $A^{\ominus} \neq \emptyset$.

Since $A^{\ominus}$ neither contains a cycle nor is empty, it can consist only of a single edge: $A^{\ominus} = \{(black, white)\}$ or $A^{\ominus} = \{(white, black)\}$. Because of the symmetry of the colors, we only consider the first case. Suppose that the color white lies on a cycle in $A^{\oplus}$ (either because of a self-loop at the white color as in 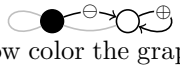 or because of a cycle involving both colors as in). We can now color the graph as follows: Color all vertices in S according to the method of Lemma 3.15 (either all of them are white or we alternate between white and black according to the distance to a fixed vertex of each component) and setup the witness function w on S. Then some vertex $v_0 \in S$ will be colored white (typically, many are white, but at least one vertex will be white). Color all vertices in $V \setminus S$ black and set the witness function to $w(v) = v_0$ for $v \in V \setminus S$. Clearly, there will be no edges between v and v_0 and, thus, the $\ominus$-arc from black to white is saturated.

Now suppose that the color white does not lie in a cycle in $A^{\oplus}$. With most cases ruled out above, the only way this can happen is when there is a $\oplus$-self-cycle at black, there is the assumed $\ominus$-arc from black to white, and possibly an $\oplus$-arc back from white to black: 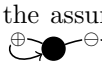 or 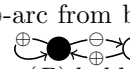 . Clearly, in the first case, where the backward $\oplus$-arc is missing, $B \notin$ SATURATION(P) holds since the vertices in S *must* be colored black and there is no way to then color the vertices in $V \setminus S$. Thus, let us now concentrate on the case 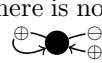 . We distinguish three cases:

1. *B consists of a single edge $\{u, v\}$ plus some isolated vertices.* Then we must have $B \notin$ SATURATION(P): We must color all isolated vertices, the vertices in $V \setminus S$, black

since there cannot be an edge from them to their witness in B and $(\mathit{black}, \mathit{white})$ is the only edge in $A^{\ominus}$. Then at least one of the two endpoints of the single edge in B (say, u) must be white, namely the endpoint that is the witness of at least one vertex in $V \setminus B$. This enforces that the other endpoint, v, is black (since $(\mathit{white}, \mathit{black}) \in A^{\oplus}$ is the only edge starting at the color white in the pattern graph). Then v cannot have a witness: The vertex u is white, so no edge in $A^{\oplus}$ can be used, nor is any of the other vertices in $V \setminus S$ white, so the edge in $A^{\ominus}$ cannot be used either.

2. *B restricted to S is a matching with at least two edges.* In this case, pick the first two edges $\{v_1, v_2\} \in E$ and $\{v_3, v_4\} \in E$ and color v_1 in white, v_2 in black, v_3 in white, and v_4 in black. Define the witness function w by $w(v_1) = v_2$, $w(v_2) = v_3$, $w(v_3) = v_4$, and $w(v_4) = v_1$. Clearly, the coloring and the witness function are correct on the vertex set $\{v_1, v_2, v_3, v_4\}$. Extend this to a coloring of all vertices as follows: All vertices of $S \setminus \{v_1, \ldots, v_4\}$ are black and their witness is the other end of the edge they are attached to, all vertices of $V \setminus S$ are black and their witness is v_1 (which is white and there is no edge in B between vertices in $V \setminus S$ and $v_1 \in S$).

3. *At least one connected component of B contains 3 or more vertices.* Let C be such a component. Consider a spanning tree T of C and let v be a leaf of this tree. Color v white and all other vertices in the component black. The witness of v is its neighbor u in the spanning tree. The witness of u is any of its neighbors other than v (such a vertex must exist since the spanning tree contains a path of length at least 2). The witnesses of all other vertices in the component is any of their neighbors in the spanning tree. Clearly, each vertex of the component is now connected by an edge in E to a black witness as required by $A^{\oplus}$. Now color all remaining vertices of S black, make any of their neighbors in B their witnesses, color all vertices of $V \setminus S$ black, and make v their witness. As in the previous case, all vertices of $V \setminus S$ now have a white witness and there is no edge between them and the witness; which is exactly what $A^{\ominus}$ requires. □

We are left with the case that the set $A^{\oplus} \cup A^{\ominus}$ contains a cycle, but neither $A^{\oplus}$ nor $A^{\ominus}$ does. This is only possible when P is either 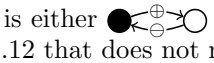 or 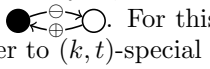 . For this special kind of cycle, there is an analogue of Fact 3.12 that does not refer to (k, t)-special graphs:

Fact 3.17 ([11, Lemma 6.7])**.** For every pattern graph P, we have $B \in \text{SATURATION}(P)$ for all B that contain a self-saturating cycle for P on which $\oplus$- and $\ominus$-arcs alternate.

Lemma 3.18. *Let $P = (\{\mathit{black}, \mathit{white}\}, A^{\oplus}, A^{\ominus})$ contain a cycle in $A^{\oplus} \cup A^{\ominus}$, but none in $A^{\oplus}$ nor in $A^{\ominus}$. Then* $\text{SATURATION}(P) \in \text{FO}$.

Proof. Let B be a basic input graph. We wish to test whether B contains a mixed self-saturating cycle for P, which must be 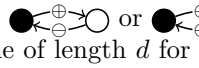 or 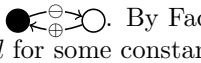 . By Fact 3.11, if such a mixed self-saturating cycle exists, there is one of length d for some constant d. (The proof in [11] yields $d = 2^{76} + 2$ for our pattern graph; but a direct argument shows that $d = 4$ suffices, fortunately.) Thus, the following formula tells us whether a mixed self-saturating cycle exists in B for P:

$$\exists a \exists b \exists c \exists d \big(E(a,b) \wedge \neg E(b,c) \wedge E(c,d) \wedge \neg E(d,a) \wedge a \neq b \wedge b \neq c \wedge c \neq d \wedge a \neq c \wedge b \neq d \wedge a \neq d \big).$$

We claim that this formula *also* tells us whether $B \in \text{SATURATION}(P)$ holds: The existence a mixed self-saturating cycle in B is a necessary condition for $B \in \text{SATURATION}(P)$ by Lemma 3.10. It is also a sufficient condition by Fact 3.17 because of the special structure of the only cycle in P. □

4 Conclusion

In the present paper we have completely classified the first-order reduction closures of prefix classes of ESO logic over directed, undirected, and basic graphs: each one of them is equal to one of the standard classes FO, L, NL, or NP. It turned out that the prefix classes for directed and undirected graphs are always the same, but often differ from the prefix classes for basic graphs. Especially interesting prefixes that mark the border between one complexity class and the next are E_1ae, E^*ae, and Eaa.

A natural question that arises is: Can we find a prefix class whose reduction closure is P? By the results of the present paper, this cannot be an ESO prefix class, unless unlikely collapses occur. However, what about prefix classes of general second-order logic? We may similarly ask whether any class other than L, NL, and the classes of the polynomial hierarchy can be characterized by a prefix class of second-order logic.

Together with the results from [6], we now have a fairly complete picture of the complexity of all ESO prefix classes over directed graphs, undirected graphs, basic graphs, and strings. Concerning arbitrary logical structures, Gottlob et al. [11] already point out that their P-NP-dichotomy for directed graphs generalizes to the collection of all finite structures over any relational vocabulary that contains a relation symbol of arity at least two; and it is not hard to see that our Theorem 1.1 also generalizes in this way (a closer look at the FO and NL upper bounds in [11] shows that they hold for arbitrary structures). The complexity of prefix classes over other special structures is, however, still open, including those of trees, infinite words, and bipartite graphs.

References

[1] Eric Allender, Michael Bauland, Neil Immerman, Henning Schnoor, and Heribert Vollmer. The complexity of satisfiability problems: Refining Schaefer's theorem. *Journal of Computer and System Sciences*, 75(4):245–254, 2009.

[2] Egon Börger, Erich Grädel, and Yuri Gurevich. *The Classical Decision Problem*. Springer-Verlag, Berlin, 1997.

[3] Julius R. Büchi. Weak second-order arithmetic and finite automata. *Zeitschrift für mathematische Logik und Grundlagen der Mathematik*, 6:66–92, 1960.

[4] Stephen A. Cook and Pierre McKenzie. Problems complete for deterministic logarithmic space. *Journal of Algorithms*, 8(5):385–394, 1987.

[5] Bruno Courcelle. The monadic second-order logic of graphs. I. Recognizable sets of finite graphs. *Information and Computation*, 85(1):12–75, 1990.

[6] Thomas Eiter, Georg Gottlob, and Yuri Gurevich. Existential second order logic over strings. *Journal of the ACM*, 47(1):77–131, 2000.

[7] Michael Elberfeld, Martin Grohe, and Till Tantau. Where first-order and monadic second-order logic coincide. In *Proceedings of the 27th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS 2012)*, pages 265–274. IEEE Computer Society, 2012.

[8] Michael Elberfeld, Andreas Jakoby, and Till Tantau. Logspace versions of the theorems of Bodlaender and Courcelle. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS 2010)*, pages 143–152, 2010.

[9] Ronald Fagin. Generalized first-order spectra and polynomial-time recognizable sets. *Complexity of Computation*, 7:43–74, 1974.

[10] Stéphane Földes and Peter L. Hammer. Split graphs. In *Proceedings of the Eighth Southeastern Conference on Combinatorics, Graph Theory and Computing*, Congressus Numerantium XIX, pages 311–315. Louisiana State Univeristy, Baton Rouge, Louisiana, 1977.

[11] Georg Gottlob, Phokion G. Kolaitis, and Thomas Schwentick. Existential second-order logic over graphs: Charting the tractability frontier. *Journal of the ACM*, 51(2):312–362, 2004.

[12] Edith Hemaspaandra, Holger Spakowski, and Mayur Thakur. Complexity of cycle length modularity problems in graphs. In *Proceedings of the 6th Latin American Symposium on Theoretical Informatics (LATIN 2004)*, volume 2976 of *Lecture Notes in Computer Science*, pages 509–518. Springer, 2004.

[13] Neil Immerman. *Descriptive Complexity Theory*. Springer-Verlag, New York, 1998.

[14] Neil Robertson and P. D. Seymour. Graph minors. II. Algorithmic aspects of tree-width. *Journal of Algorithms*, 7(3):309–322, 1986.

[15] Thomas J. Schaefer. The complexity of satisfiability problems. In *Proceedings of the 10th Symposium on Theory of Computing (STOC 1978)*, pages 216–226. ACM Press, 1978.